EXHIBIT O

Exhibit 7 - Dockser

'156 Patent

Claim Limitation (Claim 7)	Exemplary Disclosure
[**156a**] A device comprising:	Dockser discloses a "computing system," which is a device, comprising, among other things, a floating-point processor ("FPP"). Dockser also discloses a "main processor," which may in turn be part of a computing system, that comprises a floating-point processor ("FPP"). *See*, *e.g.*: "[C]omputing systems often incorporate floating-point processors." Dockser at [0001]. "The floating-point processor 100 may be implemented as part of the main processor, a coprocessor, or a separate entity connected to the main processor through a bus or other channel." Dockser at 0015. "The various illustrative logical units, blocks, modules, circuits, elements, and/or components described in connection with the embodiments disclosed herein may be implemented or performed in a floating-point processor that is part of a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic component, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein." Dockser at [0035]. "The methods or algorithms described in connection with the embodiments disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two." Dockser at [0035].
[**156b**] at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second	Dockser discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value, namely, the FPP, which contains a floating-point operator (FPO). Dockser performs multiplication at reduced precision on one or more floating-point operands to produce an output signal representing the product computed using Dockser's techniques. *See*, *e.g.*:

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
numerical value,	"Floating-point processors are specialized computing units that perform certain mathematical operations, e.g., multiplication, division, trigonometric functions, and exponential functions, at high speed." Dockser at [0001]. "The floating-point processor 100 includes a floating-point register file (FPR) 110; a floating-point controller (CTL) 130; and a floating-point mathematical operator (FPO) 140." Dockser at [0015]. "[T]he floating-point register file 110 includes several addressable register locations 115-1 (REG1), 115-2 (REG2), . . . 115-N (REGN), each configured to store an operand for a floating-point operation. The operands may include, for example, data from a memory and/or the results of previous floating-point operations. Instructions provided to the floating-point processor may be used to move the operands to and from the main memory." Dockser at [0016]. "Each register location 200 is configured to store a 32-bit binary floating-point number, in an IEEE-754 32-bit single format." Dockser at [0017]. "The subprecision select bits may also be used to remove power from logic in the floating-point operator FPO 140 that is not used when the selected subprecision is reduced." Dockser at [0018]. "The floating-point operator 140 may include one or more components configured to perform the floating-point operations. These components may include, but are not limited to, computational units such as a floating-point adder (ADD) 142 configured to execute floating-point add and subtract instructions, and a floating-point multiplier (MUL) 144 configured to execute floating-point multiply instructions. As seen in FIG. 1, each of the computational units ADD 142 and MUL 144 in the floating-point operator 140 is coupled to each other and to the floating-point register file 110 in a way as to allow operands to be transferred between the computational units, as well as between each computational unit and the floating-point register file 110. … The floating-point register file 110 may be used for storing intermediate results, as well as the results that are output from the floating-point operator 140." Dockser at [0019].

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	"[T]he floating-point operator 140 is coupled to the floating-point register file 110 so that for each instruction of a requested floating-point operation, the relevant computational unit, i.e. the adder 142 or the multiplier 144, can receive from the floating-point register file 110 one or more operands stored in one or more of the register locations." Dockser at [0023]. "Upon receiving the operands from the floating-point register file 110, one or more computational units in the floating-point operator 140 may execute the instructions of the requested floating-point operation on the received operands, at the subprecision selected by the floating-point controller 130. The output may be sent back to the floating-point register 110 for storage, as shown in FIG. 1." Dockser at [0024].

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	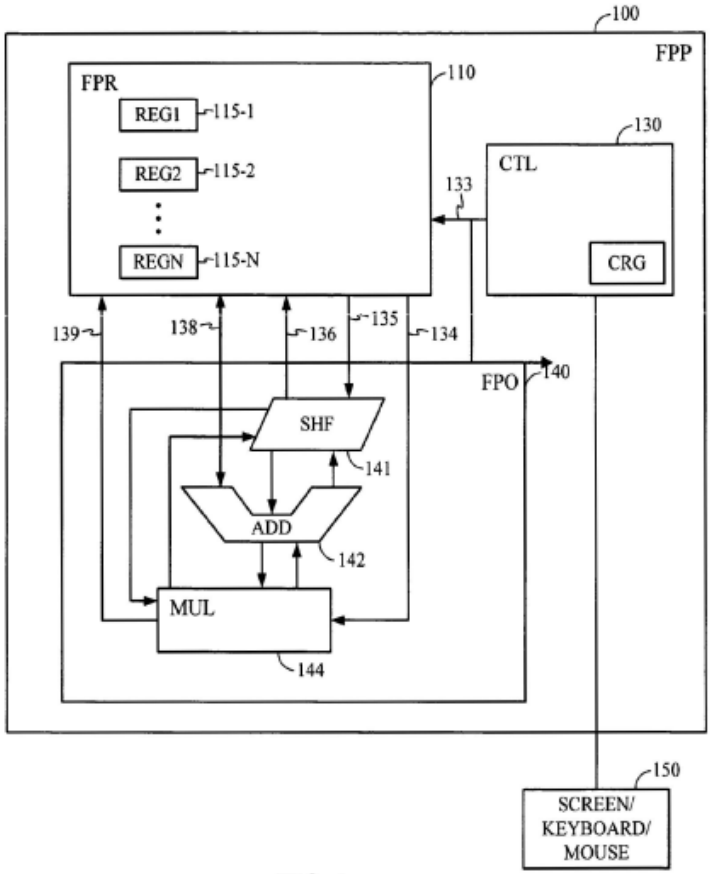 FIG. 1 Dockser, Fig. 1.
[**156c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000	Dockser discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. Specifically, as shown below, Dockser discloses removing power from portions of the FPP, including from the register file and floating-point operator in the FPP, in order to cause the FPP to perform floating-point operations at reduced precision on operands that have a dynamic range of at least 1/1,000,000 through 1,000,000. *See*, *e.g.*: "[T]he ANSI/IEEE-754 standard (commonly followed by modern computers) specifies a 32-bit single format having a 1-bit sign, an 8-bit exponent, and a 23-bit mantissa." Dockser at [0002]. "An aspect of a method of performing a floating-point operation with a floating-point processor having a precision format is disclosed. The method includes selecting a subprecision for the floating-point operation on one or more floating-point numbers, the selection of the subprecision resulting in one or more excess bits for each of the one or more floating-point numbers." Dockser at [0004]. "The floating-point processor includes a floating-point controller configured to select a subprecision for a floating-point operation on one or more floating-point numbers, the selection of the subprecision resulting in one or more excess bits for each of the one or more floating-point numbers, the floating-point controller being further configured to remove power from one or more components in the floating-point processor that would otherwise be used to store or process the one or more excess bits. The floating-point processor further includes a floating-point operator configured to perform the floating-point operation." Dockser at [0005]. "The floating-point processor includes a floating-point register having a plurality of storage elements configured to store a plurality of floating-point numbers, and a floating-point operator configured to perform a floating-point operation on the one or more of the floating-point numbers stored in the floating-point register. The floating-point processor further includes a floating-point controller configured to select a subprecision for a floating-point operation on said one or more of the floating-point numbers, the selection of the subprecision resulting in one or more excess bits for each of said one or more of the floating-point numbers, the one or

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	more excess bits being stored in one or more of the storage elements of the floating-point register, and wherein the floating-point controller is further configured to remove power from the storage elements for the one or more excess bits." Dockser at [0006]. "The floating-point processor includes a floating-point register configured to store a plurality of floating-point numbers, and a floating-point operator having logic configured to perform a floating-point operation on the one or more of the floating-point numbers stored in the floating-point register. The floating-point processor further includes a floating-point controller configured to select a subprecision for a floating-point operation on said one or more of the floating-point numbers, the selection of the subprecision resulting in one or more excess bits for each of said one or more of the floating-point numbers, and wherein the floating-point controller is further configured to remove power a portion of the logic that would otherwise be used to process the one or more excess bits." Dockser at [0007]. "The subprecision select bits may also be used to remove power from logic in the floating-point operator FPO 140 that is not used when the selected subprecision is reduced." Dockser at [0018]. "Upon receiving the operands from the floating-point register file 110, one or more computational units in the floating-point operator 140 may execute the instructions of the requested floating-point operation on the received operands, at the subprecision selected by the floating-point controller 130." Dockser at [0024]. "[A] software selectable mode may be used to reduce the precision of the floating-point operations under program control or as explained above, the instructions provided to the floating-point processor 100 may include a programmable control field containing the subprecision select bits." Dockser at [0025]. "The subprecision select bits may be used to reduce the precision of the floating-point operation. This may be achieved in a variety of ways. In one embodiment, the floating-point controller 130 may cause power to be removed from the floating-point register elements for the excess bits of the fraction that are not required to meet the precision specified by the

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	subprecision select bits. By way of example, if each location in the floating-point register file contains a 23-bit fraction, and the subprecision required for the floating-point operation is 10-bits, only the 9 commonly significant bits (MSBs) of the fraction are required; the hidden or integer bit makes the tenth. Power can be removed from the floating-point register elements for the remaining 14 fraction bits." Dockser at [0026]. 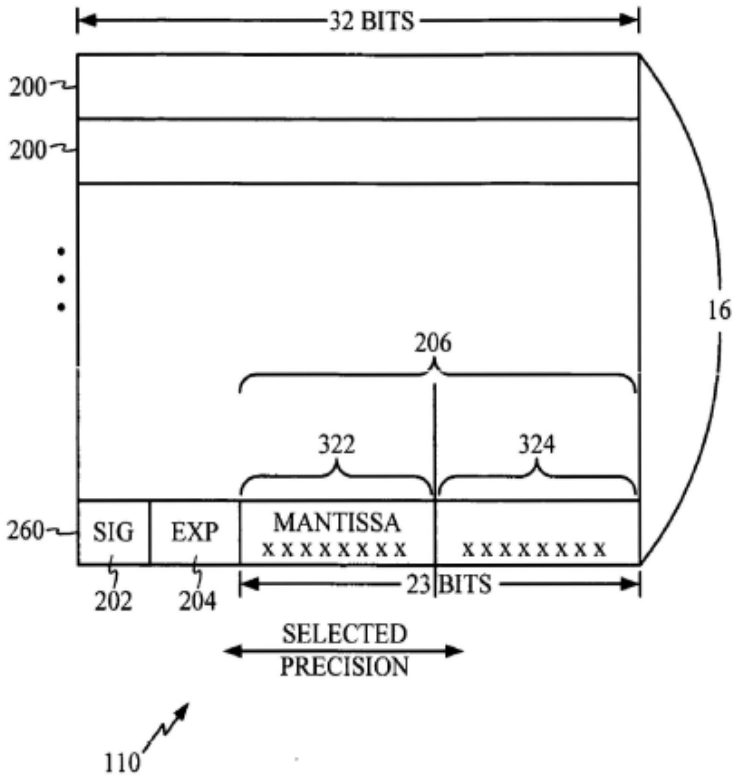 FIG. 2 Dockser, Fig. 2 "In addition, the logic in the floating-point operator 140 corresponding to the excess mantissa

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	bits do not require power." Dockser at [0027]. "In accordance with standard convention, the floating-point register stores in order the bits that make up each number, ranging from a rightmost LSB to a leftmost MSB. Each successive one of the stages moving from right to left across the FIG. 3A involves bits that have an increased significance compared to bits involved in the preceding stages." Dockser at [0028]. "In FIG. 3A, the powered bits provided to the active stages of the floating-point operator are shown as Xs, using reference numeral 322, while the unpowered bits provided to the stages with power removed are shown as circles, using reference numeral 324." Dockser at [0029]. "FIG. 3B is a conceptual diagram illustrating an example of a floating-point multiplication operation with power being selectively applied to logic in the floating-point operator." Dockser at [0030]. "As in the case of floating-point addition, floating-point multiplication is performed in a series of stages, illustrated in FIG. 3B as 410-1, . . . , 410-m." Dockser at [0031]. "In the embodiment illustrated in FIG. 3B, the selection of a desired reduced precision by the controller 130 is indicated with a line 405. As in the case of floating-point addition, described in conjunction with FIG. 3A, power may be removed from the logic used to implement the stages to the right of the line 405. Power is only applied to the stages that are actually needed to support the selected subprecision, i.e., the stages to the left of the line 405. In FIG. 3B, the bits provided to the powered on logic are shown as Xs, while the bits provided to the powered down stages are shown as circles." Dockser at [0032]. "As seen from FIG. 3B, for the first partial product 420-1, the logic for a number of bits N, shown using reference numeral 402, is unpowered. For the second partial product, the logic for N−1 bits is unpowered, and so forth. For the m-th partial product or shifted floating-point number 420-m, the logic for a number (N−m+1) of bits, shown using reference numeral 414, is unpowered." Dockser at [0033].

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	The output value 430 may be truncated to the selected subprecision, i.e. any of the bits of the output value 430 that are in less the selected precision may be truncated, to generate a truncated output number characterized by the selected precision. Alternatively, the output value 430 may be rounded off to the selected precision. In either case the output bits less significant than the selected precision may also be unpowered." Dockser at [0034]. 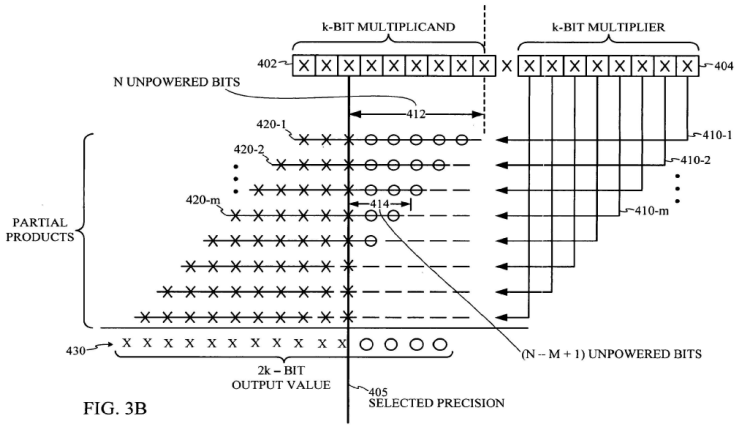 FIG. 3B Dockser, Fig. 3B. One of ordinary skill in the art would have understood Dockser to teach tying the output of un-powered components to low voltage, making their values 0, or alternatively would have understood that this was the obvious way to implement Dockser's teachings. *See*, *e.g.*, U.S. Patent No. 5,666,071 ("Hawkins") at 1:13-2:15 ("Complimentary metal-oxide-semiconductor (CMOS) has become popular when used in low power integrated circuit devices" [like Dockser's]. … "In order to retain the low power advantages of a CMOS device, it is important to ensure the gate terminal voltage not be allowed to float during periods of non-use. If the gate terminal voltage is left to float between, e.g., a power supply and ground, then that device as well as possibly other connected devices may turn on thereby causing momentary shorting of the power supply. … The gate terminal of a CMOS transistor which receives a conductor having a floating value signal will draw unacceptably large amounts of power. CMOS devices draw the least amount of power when their inputs are at the upper supply level… or at ground

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
	voltage level. If the gate terminals of those devices are allowed to float between power and ground, then significantly larger amounts of power will be consumed."); U.S. Patent Appl. Pub. No. 2003/0204760 ("Youngs") at [0005] ("Modern processors" [like Dockser's FPP] "are typically fabricated using Complementary Metal Oxide Semiconductor (CMOS) circuitry."), U.S. Patent Appl. Publ. No. 2009/0066164 ("Flynn") at [0013] ("integrated circuit design" tries "to prevent floating inputs or outputs arising when portions of the integrated circuit are powered down which could otherwise result in unpredictable or incorrect operation elsewhere"); U.S. Patent No. 5,442,577 ("Cohen") at 3:63-66, 4:51-55 (interpreting low voltage as "0" is "typical[]…[in] digital computer system[s]"); Dockser at [0029] ("[t]he carry-out C from the last powered down stage 310i is forced to zero"). *See also* Appendix to Responsive Contentions Regarding Non-Infringement and Invalidity ("Responsive Contentions") (detailing error rates associated with different mantissa sizes).
[**156d**] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit;	Dockser discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. Dockser discloses, among other things, that the precision at which the FPP operates is controlled by a processor that configures the FPP. *See*, *e.g.*, "Referring back to FIG. 1, the floating-point controller 130 may be used to select the subprecision of the floating-point operations using a control signal 133. A control register (CRG) 137 may be loaded with subprecision select bits for example transmitted in the control field of one or more instructions." Dockser at [0018]. "[A] software selectable mode may be used to reduce the precision of the floating-point operations under program control or as explained above, the instructions provided to the floating-point processor 100 may include a programmable control field containing the subprecision select bits. The subprecision select bits are written to the control register 137, which in turn controls the bit length of the mantissa for each operand during the floating-point operation. … In another embodiment of the floating-point processor 100, the subprecision selection bits may be written to the control register 137 directly from the main processor, or its operating system." Dockser at [0025].

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
[**156e**] wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine;	Dockser discloses that the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine. Dockser discloses that the main processor controlling the FPP can be a state machine. *See*, *e.g.*, "[T]he FPP "may be implemented or performed in a floating-point processor that is part of a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic component, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general-purpose processor may be a microprocessor, but in the alternative, the processor may be any conventional processor, controller, microcontroller, or ***state machine***." Dockser at [0035] (emphasis added). Alternatively, Dockser discloses that the main processor controlling the FPP is a "conventional processor," which a POSA would have understood is a CPU; alternatively, a POSA would have found it obvious to use a CPU as the main processor. *See*, *e.g.* Dockser at [0035] ("The various illustrative logical units, blocks, modules, circuits, elements, and/or components described in connection with the embodiments disclosed herein may be implemented or performed in a floating-point processor that is part of a general purpose processor … A ***general-purpose processor may be a microprocessor***, but in the alternative, the processor ***may be any conventional processor***, controller, microcontroller, or state machine."); U.S. Patent No. 6,311,282 ("Nelson") at 1:16-28 ("Portable computers are well known, as are personal 'communicators' of the type exemplified by the Motorola Envoy. Such portable computing devices are invariably battery powered. … Therefore there is known a wide range of techniques for conserving power in such battery powered devices. These power conservation methods include shutting down portions (various subsystems) of the computer when not in use, as well as putting the computer CPU (the main processor) to 'sleep' when its capabilities are not being used.").
[**156f**] and, wherein the number of LPHDR execution units in the device exceeds by at least one	Dockser discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

Exhibit 7 - Dockser

Claim Limitation (Claim 7)	Exemplary Disclosure
hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	Specifically, Dockser teaches implementing its FPP in a device that includes a combination of components, and a POSA would have understood that such a combination could include at least one hundred more FPPs (containing FPOs) than execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. *See* Dockser at [0035] ("The various illustrative logical units, blocks, modules, circuits, elements, and/or components described in connection with the embodiments disclosed herein may be implemented or performed in a floating-point processor that is part of a general purpose processor, … or other programmable logic component … , discrete gate or transistor logic, discrete hardware components, or any combination thereof … A general-purpose processor may be a microprocessor, but in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. The processor may also be implemented as a combination of computing components, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration.") To the extent Singular contends that Dockser does not disclose this limitation, the limitation would have been obvious to a person of ordinary skill in the art as any person of ordinary skill in the art would have been motivated to achieve the power-optimization described by Dockser in a system with large numbers of low-precision high-dynamic range units, such as massively parallel processing supercomputers with large numbers of processing elements configured to perform floating-point arithmetic or personal computers utilizing a single-instruction, multiple datastream architecture. *See* Responsive Contentions.

Exhibit 7 - Dockser

’273 Patent

Claim Limitation (Claim 53)	Exemplary Disclosure
[**273a**] A device:	Dockser discloses a device. *See* [**156a**].
[**273b**] comprising at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	Dockser discloses at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**273c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;	Dockser discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**].
[**273d**] wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	Dockser discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. *See* [**156f**].

Exhibit 7 - Dockser

'961 Patent

Claim Limitation (Claim 4)	Exemplary Disclosure
[**961a**] A device comprising:	Dockser discloses a device. *See* [**156a**].
[**961b**] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	Dockser discloses at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**961c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	Dockser discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**]. To the extent Singular contends that Dockser does not disclose this limitation, the limitation would have been obvious to a person of ordinary skill in the art as any person of ordinary skill in the art would have been motivated to configure Dockser's FPP to use a subprecision that meets the claims. Dockser teaches that the subprecision at which the FPP/FPO operates is a result-effective variable that affects the precision of the output of Dockser's multiplication operation, and also expressly teaches reducing precision to whatever level is necessary for a given application in order to save power. *See*, *e.g.*, Dockser at [0002] ("The precision of the floating-point processor is defined by the number of bits used to represent the mantissa. The more bits in the mantissa, the greater the precision."), [0003] ("While some applications may require these types of precision, other applications may not. For example, some graphics applications may only require a 16-bit mantissa. For these graphics applications, any accuracy

Exhibit 7 - Dockser

Claim Limitation (Claim 4)	Exemplary Disclosure
	beyond 16 bits of precision tends to result in unnecessary power consumption. … Accordingly, there is a need in the art for a floating-point processor in which the reduced precision, or subprecision, of the floating-point format is selectable."), [0014] ("In at least one embodiment of a floating- point processor, the precision for one or more floating-point operations may be reduced from that of the specified format. … By selecting the subprecision of the floating-point format, to that needed for a particular operation, thereby reducing the power consumption of the floating-point processor to support the selected subprecision, greater efficiency as well as significant power savings can be achieved."), [0024] ("one or more computational units in the floating-point operator 140 may execute the instructions of the requested floating-point operation on the received operands, at the subprecision selected by the floating-point controller 130"), [0025] ("The subprecision select bits are written to the control register 137, which in turn controls the bit length of the mantissa for each operand during the floating-point operation."), [0026] ("the floating-point controller 130 may cause power to be removed from the floating-point register elements for the excess bits of the fraction that are not required to meet the precision specified by the subprecision select bits"), [0027] ("[T]he logic in the floating-point operator 140 corresponding to the excess mantissa bits do not require power. Thus, power savings may be achieved by removing power to the logic in the floating-point operator 140 that remains unused as a result of the subprecision selected."). To the extent that reducing precision to a level that meets this limitation is not obvious in view of Dockser alone, this would have been obvious in view of prior art that expressly taught lower precision levels than those explicitly mentioned in Dockser. *See*, *e.g.*, Tong, 279 & Fig. 6; *see also* Responsive Contentions.
[**961d**] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.	Dockser discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See* [**156d**].

Exhibit 7 - Dockser

Claim Limitation (Claim 13)	Exemplary Disclosure
[**961e**] A device comprising:	Dockser discloses a device. *See* [**156a**].
[**961f**] a plurality of components comprising:	Dockser's "computing system," which meets the "device" of limitation [**961e**], contains a plurality of components. *See*, *e.g.*, [**156b**] + **[156d**].
[**961g**] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	Dockser discloses at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**961h**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of	Dockser discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**]. To the extent Singular contends that Dockser does not disclose this limitation, the limitation would have been obvious to a person of ordinary skill in the art as any person of ordinary skill in the art would have been motivated to configure Dockser's FPP to use a subprecision that meets the claims. Dockser teaches that the subprecision at which the FPP/FPO operates is a result-effective variable that affects the precision of the output of Dockser's multiplication operation, and also expressly teaches reducing precision to whatever level is necessary for a given application in order to save power. *See*, *e.g.*, Dockser at [0002] ("The precision of the floating-point processor is defined by the number of bits used to represent the mantissa. The more bits in the mantissa, the greater the precision."), [0003] ("While some applications may

Exhibit 7 - Dockser

Claim Limitation (Claim 13)	Exemplary Disclosure
the first operation on the numerical values of that same input.	require these types of precision, other applications may not. For example, some graphics applications may only require a 16-bit mantissa. For these graphics applications, any accuracy beyond 16 bits of precision tends to result in unnecessary power consumption. … Accordingly, there is a need in the art for a floating-point processor in which the reduced precision, or subprecision, of the floating-point format is selectable."), [0014] ("In at least one embodiment of a floating- point processor, the precision for one or more floating-point operations may be reduced from that of the specified format. … By selecting the subprecision of the floating-point format, to that needed for a particular operation, thereby reducing the power consumption of the floating-point processor to support the selected subprecision, greater efficiency as well as significant power savings can be achieved."), [0024] ("one or more computational units in the floating-point operator 140 may execute the instructions of the requested floating-point operation on the received operands, at the subprecision selected by the floating-point controller 130"), [0025] ("The subprecision select bits are written to the control register 137, which in turn controls the bit length of the mantissa for each operand during the floating-point operation."), [0026] ("the floating-point controller 130 may cause power to be removed from the floating-point register elements for the excess bits of the fraction that are not required to meet the precision specified by the subprecision select bits"), [0027] ("[T]he logic in the floating-point operator 140 corresponding to the excess mantissa bits do not require power. Thus, power savings may be achieved by removing power to the logic in the floating-point operator 140 that remains unused as a result of the subprecision selected."). To the extent that reducing precision to a level that meets this limitation is not obvious in view of Dockser alone, this would have been obvious in view of prior art that expressly taught lower precision levels those explicitly mentioned in Dockser. *See*, *e.g.*, Tong, 279 & Fig. 6; *see also* Responsive Contentions.